

What is claimed is:

1. A high performance relational database management system, leveraging the functionality of a high speed communications network, comprising the steps of:

(i) receiving collected data objects from at least one data collection node using at least one performance monitoring computer whereby a distributed database is created;

(ii) partitioning the distributed database into data hunks using a histogram routine running on at least one performance monitoring server computer;

(iii) importing the data hunks into a plurality of delegated database engine instances located on at least one performance monitoring server computer so as to parallel process the data hunks whereby processed data is generated; and

(iv) accessing the processed data using at least one performance client computer to monitor data object performance.

2. The system according to claim 1, wherein at least one database engine instance is located on the performance monitor server computers on a ratio of one engine instance to one central processing unit whereby the total number of engine instances is at least two so as to enable the parallel processing of the distributed database.

3. The system according to claim 2, wherein at least one database engine instance is used to maintain a versioned master vector table.

4. The system according to claim 3, wherein the versioned master vector table generates a histogram routine used to facilitate the partitioning of the distributed database.

5. The system according to claim 4, wherein the histogram routine comprises the steps of:

(i) dividing the total number of active object identifiers by the desired number of partitions so as to establish the optimum number of objects per partition;

(ii) generating an  $n$  point histogram of desired granularity from the active indices; and

(iii) summing adjacent histogram routine generated values until a target partition size is reached but not exceeded.

6. The system according to claim 1, wherein the performance monitor server comprises an application programming interface compliant with a standard relational database query language.

7. A high performance relational database management system, leveraging the functionality of a high speed communications network, comprising:

(i) at least one performance monitor server computer connected to the network for receiving network management data objects from at least one data collection node device whereby a distributed database is created;

(ii) a histogram routine running on the performance monitoring server computers for partitioning the distributed database into data hunks;

(iii) at least two database engine instances running on the performance monitoring server computers so as to parallel process the data hunks whereby processed data is generated; and

(iv) at least one performance monitor client computer connected to the network for accessing the processed data whereby data object performance is monitored.

8. The system according to claim 7, wherein at least one database engine instance is located on the performance monitoring server computers on a ratio of one engine instance to one central processing unit whereby the total number of engine instances for the system is at least two so as to enable the parallel processing of the distributed database.

13

9. The system according to claim 8, wherein at least one database engine instance is used to maintain a versioned master vector table.

5 10. The system according to claim 9, wherein the versioned master vector table generates a histogram routine used to facilitate the partitioning of the distributed database.

11. The system according to claim 10, wherein the histogram routine comprises the steps of:

- 10 (i) dividing the total number of active object identifiers by the desired number of partitions so as to establish the optimum number of objects per partition;
- (ii) generating an  $n$  point histogram of desired granularity from the active indices; and
- (iii) summing adjacent histogram routine generated values until a target partition size is reached but not exceeded.

15 12. The system according to claim 7, wherein the performance monitor server comprises an application programming interface compliant with a standard relational database query language.

20 13. The system according to claim 7, wherein at least one performance monitor client computer is connected to the network so as to communicate remotely with the performance monitor server computers.

25 14. A storage medium readable by an install server computer in a high performance relational database management system including the install server, leveraging the functionality of a high speed communications network, the storage medium encoding a computer process comprising:

14

(i) a processing portion for receiving collected data objects from at least one data collection node using at least one performance monitoring computer whereby a distributed database is created;

(ii) a processing portion for partitioning the distributed database into data hunks using a histogram routine running on at least one performance monitoring server computer;

(iii) a processing portion for importing the data hunks into a plurality of delegated database engine instances located on at least one performance monitoring server computer so as to parallel process the data hunks whereby processed data is generated; and

(iv) a processing portion for accessing the processed data using at least one performance client computer to monitor data object performance.

15. The system according to claim 14, wherein at least one database engine instance is located on the data processor server computers on a ratio of one engine instance to one central processing unit whereby the total number of engine instances is at least two so as to enable the parallel processing of the distributed database.

16. The system according to claim 15, wherein one of the database engine instances is designated as a prime database engine instance used to maintain a versioned master vector table.

17. The system according to claim 16, wherein the versioned master vector table generates a histogram routine used to facilitate the partitioning of the distributed database.

18. The system according to claim 14, wherein the histogram routine comprises the steps of:

15

(i) dividing the total number of active object identifiers by the desired number of partitions so as to establish the optimum number of objects per partition;

(ii) generating an n point histogram of desired granularity from the active indices;  
and

5 (iii) summing adjacent histogram routine generated values until a target partition size is reached but not exceeded.

19. The system according to claim 14, wherein the performance monitor server comprises an application programming interface compliant with a standard relational database query language.

10

10326 0349000